



Artificial Intelligence: Everything Old is New Again

Henry Kautz

Division Director, Information & Intelligent Systems (IIS)

Computer and Information Science & Engineering (CISE)

National Science Foundation

A.I. is in a 'golden age' and solving problems that were once sci-fi.

– Jeff Bezo, CEO Amazon


AI is the new electricity.

– Andrew Ng, Founder Coursera, Chief Scientist Baidu

AI is more profound than electricity or fire.

– Sundar Pichai, Google CEO


Computer Vision, circa 2005



(a) Left image: 384x288, 15 labels



(b) Ground truth



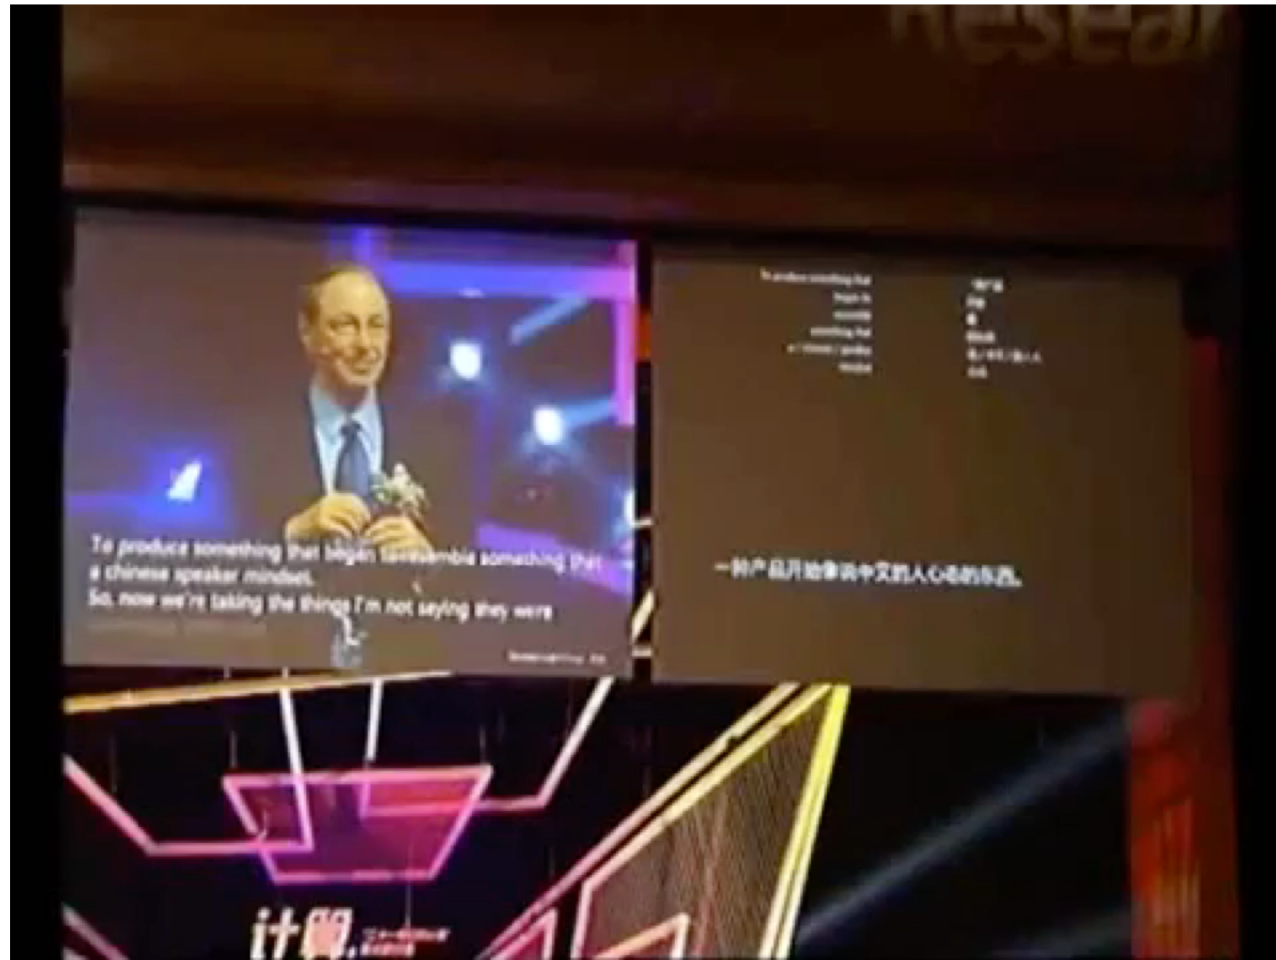
(c) Processed image

P. Felzenszwalb, D. Huttenlocher (2006), Efficient Belief Propagation for Early Vision, International Journal of Computer Vision, Vol. 70, No. 1, October 2006

Computer Vision, circa 2016




Natural Language Translation (2016)




Multilayer Artificial Neural Networks AKA Deep Learning

- Initial ideas from 1940's
- Core technical developments in 1980's
- Today's hardware 10,000 times faster – repurposing video game graphics hardware!
- Moravec's Hardware Hypothesis: *Need brain compute power*
- Also required: BIG Data!




Reflexive Intelligence

- Deep Learning excels at tasks requiring reflexive reasoning
- Near instantaneous recognition, generation, and translation of patterns
 - Perception
 - Locomotion / manipulation
 - Flight or fight decision making
- Kahneman: System I



Deliberative Intelligence

- Often overlooked in the excitement over deep learning is decades of steady progress on deliberative reasoning
 - Deduction
 - Heuristic search
 - *Sound* probabilistic reasoning
- Combinatorial optimization
 - Kahneman System 2
- Not just brute force enumeration!
 - Randomized algorithms
 - Caching partial solutions



1997: IBM's Deep Blue defeats Kasparov



Mathematical Discovery

Consider a sequence of 1s and -1s, e.g.:

-1, 1, 1, -1, 1, 1, -1, 1, -1, -1 ...
 1 -1 1 1 -1 ...
 1 1 -1 ...
 -1 1 ...

Erdos Discrepancy Conjecture:
 For all N exists some
 sequence with all sums
 between -2 and $+2$

and look at the sum of the sequence and its subsequences

-1 + 1 = 0
 -1 + 1 + 1 = 1
 -1 + 1 + 1 + -1 = 0
 -1 + 1 + 1 + -1 + 1 = 1
 -1 + 1 + 1 + -1 + 1 + 1 = 2
 -1 + 1 + 1 + -1 + 1 + 1 + -1 = 1
 -1 + 1 + 1 + -1 + 1 + 1 + -1 + 1 = 2
 -1 + 1 + 1 + -1 + 1 + 1 + -1 + 1 + -1 = 1

and "skip by 1"

1 + -1 = 0
 1 + -1 + 1 = 1
 1 + -1 + 1 + 1 = 2

and "skip by 2"

1 + 1 = 2
 1 + 1 + -1 = 1

We now know (2015): there exists a sequence of 1160 +1s and -1s such
 that sums of all subsequences *never* < -2 or $> +2$.

Superhuman Deliberative Intelligence

- Result was obtained with a *general* reasoning program - a Boolean Satisfiability or SAT solver (2015)
 - 37,462 Boolean variables
 - 161,644 constraints
- Proof of non-existence of discrepancy 2 sequence found in about 10 hours on a laptop!
 - Proof: a billion small inference steps
 - Independently verified by a simple (50 line) proof checking program
- No human could create or understand this kind of proof
- But: we can be trust of the result because of the verifier

Deliberative AI in Action

- Superhuman “chess-like” reasoning prowess solves hard problems in planning, resource optimization, and design
- NSF Expeditions Award: Computational Sustainability

Materials Discovery




Photo: John Gregoire (JCAP/Caltech)

What: Rapid characterization of crystal structures from high-throughput X-ray diffraction experiments.

Why: Identify new materials for fuel cells, energy storage, and solar fuel generation.

How: Pattern decomposition, constraint and probabilistic reasoning, crowdsourcing.

Smart Grid

What: Power grid modeling, control, and energy storage.

Why: Managing the power system with increasing use of renewable sources of electricity.

How: Stochastic optimization, sequential decision making, pattern decomposition.



Photo: DOE

Big Data for Africa



Photo: Frank Annor (TAHMO)

What: Deploy 20,000 low-cost weather stations across Africa.

Why: Improve weather predictions, which is directly related food security.

How: Optimal placement, bayesian networks, multi-scale probabilistic modeling.

Landscape-Scale Conservation

What: Socio-ecological corridor in the Ecuadorian Andes.

Why: Protect endangered Andean bear and other species in a significant

biodiversity hotspot, while improving livelihoods of local communities.

How: Spatial capture-recapture, stochastic optimization, spatio-temporal modeling.




Photo: Santiago Molina

Reflexive versus Deliberative Intelligence

- Some tasks that were thought to require deliberation can be solved reflexively by deep learning, *if justification of result not required*
 - First-pass mass scanning of medical images
 - Credit risk scoring
- Many deep learning researchers believe deep learning can eventually encompass deliberative as well as reflexive reasoning
 - Neural Turing Machine (2015) – in *principle* can learn to do anything, in *practice* can learn to sort small numbers
- Hybrid artificial intelligence: Reflexive + Deliberative Reasoning
 - Deep Learning + Combinatorial Optimization

Alpha Go / Alpha Zero (2016/2017)

- Deep neural network learns “evaluation function” – reflexive estimate of value of board position
- Stochastic tree search (kind of combinatorial optimization) uses evaluation function to choose play
- Alpha Zero: generalizes to 2 player board games of perfect information




Still Needed: *Comprehensive* Intelligence



- AKA Commonsense
- Needed to deal with unforeseen cases, not in training data
- Example: Streetsweeper driving slowly on left edge of highway
 - Human drivers easily avoid
 - Tesla on autopilot crashes into it!





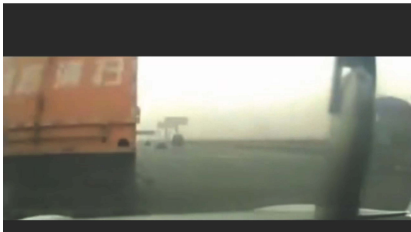
Human versus Tesla



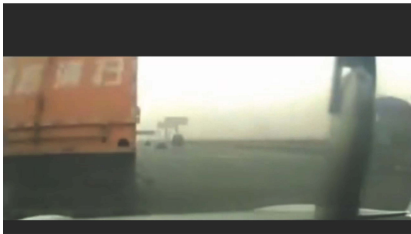
Human versus Tesla

	No one else is driving in the left lane. Do they see trouble ahead?	Wheeeee! No traffic in my lane!
		

Human versus Tesla

	No one else is driving in the left lane. Do they see trouble ahead?	Wheeeee! No traffic in my lane!
	What the heck is that orange thing? I don't know, but I better avoid it!	I don't see any cars or pedestrians!
		

Human versus Tesla

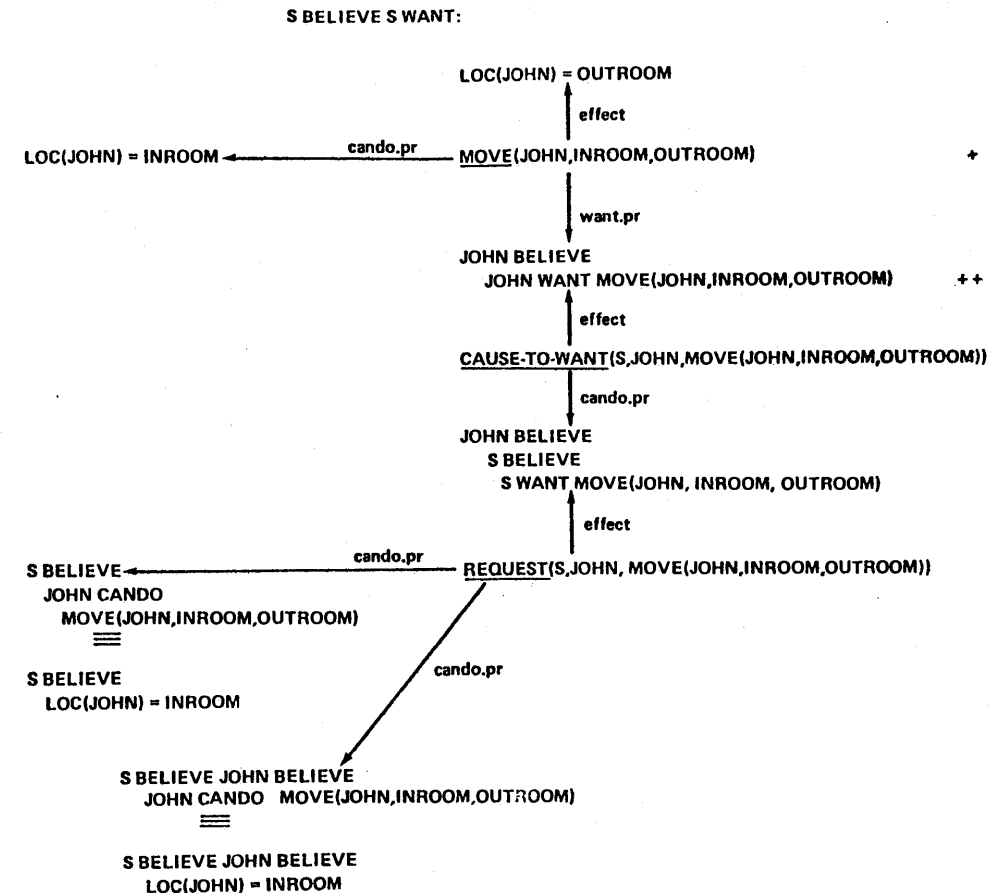
	No one else is driving in the left lane. Do they see trouble ahead?	Wheeeee! No traffic in my lane!
	What the heck is that orange thing? I don't know, but I better avoid it!	I don't see any cars or pedestrians!
	A streetsweeper on an expressway?! What a crazy thing!	[DEAD]

Natural Language (Non) Understanding

- Today's AI system may recognize and even translate natural language, but they do not *understand* it
- Fail to solve simple pronoun resolution problems that require understanding
 - Winograd Schema Challenge (Morgenstern et al. 2016)
 - The city councilmen refused the demonstrators a permit because **they feared** violence.
 - The city councilmen refused the demonstrators a permit because **they advocated** violence.


Discourse Understanding

- Back in the 1980's, researchers formalized natural language discourse as a process of **planning and plan recognition** (Cohen 1980; Allen 1983)
- Key idea: **utterances are actions** described by preconditions and effects
- Linguistic tasks such as pronoun resolution are **inferring the most likely parameters** to the underlying discourse plans
- Framework **naturally handles multiple utterance** exchanges – not just single question / answer pairs
- **Research challenge**: Employ this framework with modern methods for learning, reasoning, and speech recognition!



Representation Learning

- Deep learning from raw features (e.g. pixels) must abstract the data in some way
- The internal state of a deep network is thus in some sense a **representation** of the input
- Key question: is the representation **interpretable**?
- Not by default – but it may be possible to design neural net architectures such that it is!
- **Detangled representations**
(Siddharth et al. 2017)



What this Means for NSF AI Strategy

- While industry is investing heavily in AI applications, federal support is still crucial for fundamental research
 - Understanding how, why, and when deep learning works
 - Hybrid reasoning
 - General intelligence / commonsense
 - AI applications for social good (i.e. non-commercial)
- *Reducing deliberative tasks to reflexive ones can lead to ethical quandaries*
 - Credit risk scoring
 - Prison sentencing
 - Use of lethal force in autonomous weapons


NSF CISE Dear Colleague Letter (Nov 2, 2018)

Fairness, Ethics, Accountability, and Transparency: Enabling Breakthrough Research to Expand Inclusivity in Computer and Information Science and Engineering Research

- **Fairness:** How do we prevent socially undesirable bias in AI algorithms?
- **Ethics:** What are ethical concerns when creating AI systems?
- **Accountability:** How is responsibility for the outcomes of an AI system shared between inventors, implementors, and end users?
- **Transparency:** How can and when must the inner workings of an AI system be made understandable to users?

NSF AI Programs Featuring FEAT

- NSF/CISE has been supporting fundamental AI search for 50 years
- Now: numerous AI-focused interdisciplinary funding programs that emphasize the social impacts of AI



**The Future of Work at
the Human-Technology
Frontier**

The Future of Work at the Human-Technology Frontier

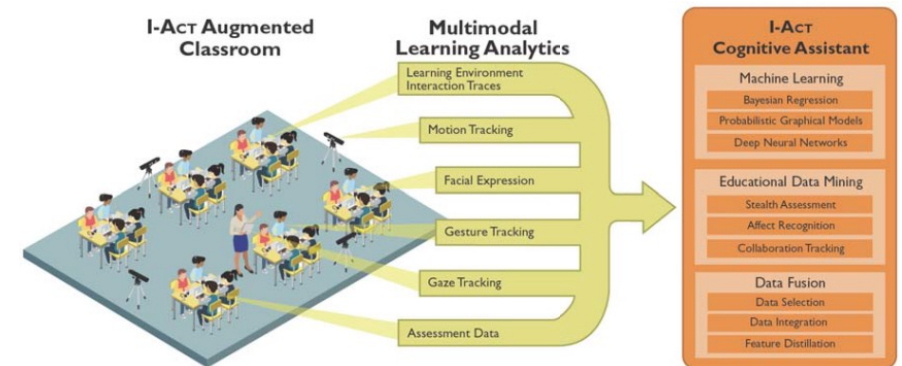
2018:

- Foundations for Augmenting Human Cognition
- Embodied Intelligent Cognitive Assistants



• 2019:

- Expanded solicitation coming out in a few days!



Development of AI Roadmaps

- Artificial Intelligence Roadmap
 - Bottom-up, driven by research community
 - Computing Community Consortium with support from NSF
 - Steering Committee from academia and industry
 - Workshops
 - November 2018 – Integrated Intelligence
 - December 2018 – Interaction
 - Winter 2019 – Learning & Robotics



CRA
Computing Research
Association



Involving the Community

- W1: Integrated Intelligence

- Chairs: Marie desJardins and Ken Forbus
- Understanding the mind
- Composing intelligent capabilities
- Open repositories of knowledge

- W2: Interaction

- Chairs: Kathy McKeown and Dan Weld
- Interactions that matter
- Trust and responsibility
- People interacting online

- W3: Learning and Robotics

- Chairs: Tom Dietterich and Fei-Fei Li
- Deeper learning
- Integrating statistical learning and symbolic representations
- Diversified learning modalities

- AAI Town Hall

- Email your ideas to cccinfo@cra.org

National AI Research & Development Strategic Plan

- **Plan for AI investment in R&D across all federal agencies**
 - Chaired by NSF and IARPA, 40 participating agencies
 - Spring 1019: R&D Plan plus Implementation Report
 - Request for Information comments received Oct 26 from individuals, universities, major high-tech companies
- **Responses stress**
 - Importance of synergy between federal and industry investments
 - Importance of supporting work in FEAT
 - US at high risk of losing leadership in AI to China

Citations

- Hans P. Moravec, Robot: Mere Machine to Transcendent Mind, Oxford University Press, 2000.
- Daniel Kahneman, Thinking, Fast and Slow, Farrar, Straus and Giroux, 2011.
- Boris Konev and Alexei Lisitsa, A SAT Attack on the Erdos Discrepancy Conjecture, in Theory and Applications of Satisfiability Testing (SAT 2014), Lecture Notes in Computer Science, vol. 8561, Springer, Cham, 2014.
- Alex Graves, Greg Wayne, and Ivo Danihelka, Neural Turing Machines, CoRR abs/1410.5401, 2014.
- David Silver, ..., Demis Hassabis, Mastering the Game of Go with Deep Neural Networks and Tree Search, Nature, vol. 529, 28 January 2016.
- Leora Morgenstern, Ernest Davis, and Charles Ortiz, Planning, Executing, and Evaluating the Winograd Schema Challenge, AI Magazine, Spring 2016.
- Philip R. Cohen and C. Raymond Perrault, Elements of a Plan-based Theory of Speech Acts, Cognitive Science, 3(3), 1979.
- J. F. Allen, J. F. and C. R. Perrault, Analyzing Intention in Utterances, Artificial Intelligence 15(3), 1980.
- N. Siddharth, ..., Philip Torr, Learning Disentangled Representations with Semi-Supervised Deep Generative Models, NIPS 2017.